

IN THE CLAIMS:

1 1. (Previously Presented) In a file server having a storage operating system, a method for
2 managing storage of data in a plurality of storage devices, each storage device having a
3 plurality of blocks for storing data, comprising:

4 generating block layout information in a file system layer of the storage operating
5 system by determining which blocks within the plurality of blocks are allocated for
6 storing data and which are unallocated;

7 transferring the block layout information from the file system layer to a RAID
8 layer of the storage operating system; and

9 responsive to the block layout information, controlling the execution of I/O
10 operations at the RAID layer by identifying a plurality of contiguous blocks on a single
11 storage device within the plurality of blocks for use by each I/O operation so as to
12 substantially maximize chain lengths of reads for calculation of parity;

13 selecting a parity subtraction method or a recalculation method for parity
14 calculation based on the method that requires a fewest number of read operations to
15 compute parity for the I/O operations; and

16 responsive to the block layout information and the parity calculation method
17 selected, identifying the blocks within the plurality of blocks for use by the I/O
18 operations.

1 2. (Previously Presented) A method for managing storage of data in a plurality of storage
2 devices, each comprising a plurality of storage blocks, comprising:

3 generating block layout information; and

4 in response to the block layout information, controlling execution of I/O operations
5 by identifying a plurality of contiguous storage blocks on a single storage device for use

6 by each I/O operation so as to substantially minimize a number of read operations needed
7 for calculation of error correction parameters across a stripe.

1 3. (Original) The method of claim 2 wherein the calculation of error correction
2 parameters comprises the calculation of parity.

1 4. (Original) The method of claim 3 wherein the calculation of parity comprises selecting
2 a parity calculation operation from a group consisting of a subtraction method and a
3 parity re-calculation method.

1 5. (Previously Presented) The method of claim 2 wherein the identification of storage
2 blocks for use in the I/O operation substantially maximizes a chain length by
3 substantially maximizing the number of blocks having a contiguous physical layout on
4 the storage device.

1 6. (Previously Presented) The method of claim 2, further comprising:
2 identifying storage blocks for use in the I/O operation so as to substantially
3 maximize the chain length by substantially maximizing the number of blocks having
4 sequential volume block numbers (VBNs) associated with the storage blocks.

1 7. (Previously Presented) The method of claim 2, further comprising:
2 identifying storage blocks for use in the I/O operation so as to substantially
3 maximize the chain length by substantially maximizing locality of the blocks of the
4 storage device.

1 8. (Previously Presented) The method of claim 2 wherein the controlling execution step
2 comprises:

3 examining blocks to which data is to be written prior to write operations;
4 selecting one of a plurality of parity calculation methodologies including,
5 a first methodology comprising minimizing a number of blocks read, and
6 a second methodology comprising maximizing chain lengths of blocks read for
7 parity calculation.

1 9. (Previously Presented) The method of claim 8, wherein the controlling execution step
2 further comprises:

3 implementing selection of the parity calculation methodology responsive to the
4 block layout information; and

5 wherein, if the selection constitutes substantially minimizing the number of blocks
6 read,

7 determining on a stripe-by-stripe basis whether to calculate parity based on a
8 subtraction method or a recalculation method,

9 performing any appropriate read operations to support the method selected, and

10 calculating parity responsive to the blocks read and the data to be written; and

11 wherein, if the selection constitutes substantially maximizing chain lengths of blocks
12 read,

13 deciding which storage blocks to read to substantially maximize chain length while
14 substantially minimizing the number of storage blocks read to support either a subtraction
15 method or a recalculation method,

16 performing read operations on the blocks to be read, and

17 calculating parity responsive to the blocks read and the data to be written.

- 1 10. (Original) The method of claim 2, wherein the identification of storage blocks is
- 2 based at least in part on an available resource.

- 1 11. (Original) The method of claim 2 further comprising transmitting the block layout
- 2 information from a file system layer to a RAID layer.

- 1 12. (Original) The method of claim 2 wherein the generating step further comprises:
 - 2 making a first determination as to whether a storage block is unallocated;
 - 3 making a second determination as to a current implementation of the plurality of
 - 4 storage devices; and
 - 5 generating the block layout information based at least in part on the first and the
 - 6 second determinations.

- 1 13. (Original) The method of claim 2, wherein the I/O operation is one of a plurality of
- 2 I/O operations and one of the plurality of I/O operations is a read operation.

- 1 14. (Original) The method of claim 2, wherein the chain length is a chain length of a read
- 2 operation for calculation of parity.

- 1 15. (Original) The method of claim 2, wherein the chain length is a chain length for a
- 2 write operation for the data.

- 1 16. (Previously Presented) A method for managing storage of data in a storage system
- 2 comprising:

3 maintaining a plurality of storage devices each having a plurality of storage
4 blocks;

5 writing data to predetermined storage blocks across a plurality of stripes and to
6 predetermined contiguous storage blocks within each storage device so as to substantially
7 maximize chain length of storage blocks within each storage device and minimizing a
8 number of read operations for the calculation of error correction parameters across each
9 stripe of the plurality of stripes by selecting a parity subtraction method or a recalculation
10 method for parity calculation based on the method that requires the fewest number of
11 read operations to compute parity.

17. - 38. (Cancelled)

1 39. (Previously Presented) A storage system comprising:

2 a plurality of storage devices each having a plurality of storage blocks; and
3 a storage manager in communication with the plurality of storage devices, the
4 storage manager writing data to predetermined storage blocks across a plurality of stripes
5 and to predetermined storage blocks within each storage device so as to substantially
6 maximize chain length of storage blocks by selecting as many as contiguous storage
7 blocks within a single storage device while substantially minimizing a number of read
8 operations required for calculation of error correction parameters across each stripe of the
9 plurality of stripes by selecting a parity subtraction method or a recalculation method for
10 parity calculation based on the method that requires the fewest number of read operations
11 to compute parity.

1 40. (Previously Presented) A system for managing the storage of data, the system
2 comprising:

3 a plurality of storage devices each having a plurality of storage blocks;

4 a storage device manager in communication with the plurality of storage blocks;
5 a block layout information generator in communication with the storage device
6 manager and the plurality of storage blocks; and
7 an error correction parameter calculator in communication with the plurality of
8 storage blocks and the storage device manager,
9 wherein the storage device manager, in response to the block layout information
10 from the block layout information generator, controls the execution of an I/O operation
11 by identifying a plurality of contiguous storage blocks on a single storage device for use
12 by the I/O operation so as to substantially maximize chain length within the storage
13 device while substantially minimizing the number of read operations required for
14 calculation by the error correction parameter calculator of error correction parameters
15 across a stripe by selecting a parity subtraction method or a recalculation method for
16 parity calculation based on the method that requires the fewest number of read operations
17 to compute parity.

1 41 - 44. (Cancelled)

1 45. (Previously Presented) A method for managing storage of data by a server,
2 comprising:
3 receiving a request to write data to a plurality of storage devices;
4 generating block layout information to determine which blocks within a plurality
5 of blocks located in the plurality of storage devices are allocated for storing data and
6 which are unallocated;
7 identifying blocks within the plurality of blocks for use by a set of I/O operations
8 to store the data;
9 determining the number of read operations needed to compute parity for the data
10 by computing parity using a subtraction method of computing parity;

11 determining the number of read operations needed to compute parity for the data
12 by computing parity using a recalculation method of computing parity;

13 choosing either the subtraction method of computing parity or the recalculation
14 method of computing parity by determining which of these two methods requires the
15 fewer number of read operations, and choosing the method requiring the fewer number of
16 read operations; and

17 writing the data to identified blocks, and computing parity for the data using the
18 chosen method of computing parity.

1 46. (Previously Presented) The method of claim 45, further comprising:

2 choosing to either firstly maximize chain lengths of read operations for
3 calculation of parity or secondly to place the data with a high degree of locality in the
4 plurality of storage devices, by choosing the method which requires the fewest number of
5 read operations in computing parity for the data.

1 47. (Currently amended) A method for managing storage of data by a server, comprising:

2 receiving a request to write data to a plurality of storage devices;

3 generating block layout information to determine which blocks within a plurality
4 of blocks located in the plurality of storage devices are allocated for storing data and
5 which are unallocated;

6 identifying blocks within the plurality of blocks for use by a set of I/O operations
7 to store the data; and

8 selecting whether to substantially minimize the number of read blocks or to
9 substantially maximize chain lengths of read blocks, and

10 implementing the selection, during the writing of the data to the plurality of
11 storage devices, responsive to the block layout information, and responsive to whether

12 substantially minimizing the number of read blocks or substantially maximizing chain
13 lengths of read blocks requires fewer number of read operations.

1 48. (Previously Presented) The method of claim 47, further comprising:

2 in response to selecting to substantially minimize the number of read blocks,
3 determining whether to calculate parity based on the subtraction method or the
4 recalculation method by determining which method requires the fewer number of read
5 operations, and selecting the method which requires the fewer number of read operations,
6 and

7 performing the write operation and calculating the parity using the parity
8 calculation method requiring the fewer number of read operations.

1 49. (Previously Presented) The method of claim 48, further comprising:

2 in response to selecting to substantially maximize chain lengths of read blocks,
3 deciding which storage blocks to read to substantially maximize chain length while
4 minimizing the number of storage blocks read to support either the subtraction method or
5 the recalculation method of parity calculation; and

6 performing the write operation and calculating the parity using the parity
7 calculation method requiring the fewer number of read operations.

1 50. (Previously Presented) A method for managing storage of data by a server,
2 comprising:

3 receiving a request to write data to a plurality of storage devices;

4 generating block layout information to determine which blocks within a plurality
5 of blocks located in the plurality of storage devices are allocated for storing data and
6 which are unallocated;

7 identifying blocks within the plurality of blocks for use by a set of I/O operations
8 to store the data;

9 testing to either maximize chain lengths of read operations for calculation of
10 parity, or to place the data with a high degree of locality in the plurality of storage
11 devices, the testing having the steps,

12 determining, for both maximizing chain length and placing the data with a high
13 degree of locality, the number of read operations needed to compute parity for the data,
14 by computing parity using both the subtraction method of computing parity and the
15 recalculation method of computing parity;

16 firstly choosing to either maximize chain lengths of read operations for
17 calculation of parity or to place the data with a high degree of locality in the plurality of
18 storage devices, and after this first choice, secondly choosing either the subtraction
19 method of computing parity or the recalculation method of computing parity by
20 determining which of these methods requires the fewest number of read operations,

21 choosing the method requiring the fewest number of read operations of computing
22 parity of the data; and

23 writing the data to identified blocks, and computing parity for the data using the
24 chosen method of computing parity.

1 51. (Previously Presented) A computer readable media, comprising:

2 said computer readable media containing instructions for execution on a processor
3 for a method of managing storage of data in a plurality of storage devices, each storage
4 device having a plurality of blocks for storing data, the method having,

5 generating block layout information; and

6 in response to the block layout information, controlling the execution of an I/O
7 operation by identifying a plurality of contiguous storage blocks on a single storage
8 device for use by the I/O operation so as to substantially minimize the number of read

9 operations needed for calculation of error correction parameters across a stripe by
10 selecting a parity subtraction method or a recalculation method for parity calculation
11 based on the method that requires the fewest number of read operations to compute
12 parity.